Supporting Reasoning with Different Types of Evidence in Intelligence Analysis

Alice Toniolo
Timothy J. Norman
Dept. of Computing Science
University of Aberdeen, UK

Anthony Etuk
Federico Cerutti
Dept. of Computing Science
University of Aberdeen, UK

Robin Wentao Ouyang
Mani Srivastava
University of California
Los Angeles, CA, USA

Nir Oren
Dept. of Computing Science
University of Aberdeen, UK

Timothy Dropps
John A. Allen
Honeywell, USA

Paul Sullivan
INTELPOINT Incorporated
Pennsylvania, USA

ABSTRACT

The aim of intelligence analysis is to make sense of information that is often conflicting or incomplete, and to weigh competing hypotheses that may explain a situation. This imposes a high cognitive load on analysts, and there are few automated tools to aid them in their task. In this paper, we present an agent-based tool to help analysts in acquiring, evaluating and interpreting information in collaboration with others. Agents assist analysts in reasoning with different types of evidence to identify what happened and why, what is credible, and how to obtain further evidence. Argumentation schemes lie at the heart of the tool, and sensemaking agents assist analysts in structuring evidence and identifying plausible hypotheses. A crowdsourcing agent is used to reason about structured information explicitly obtained from groups of contributors, and provenance is used to assess the credibility of hypotheses based on the origins of the supporting information.

Categories and Subject Descriptors

H.4 [**Inform. Systems Applications**]: Decision support

Keywords

Innovative Applications; Aerospace and Defense; Argumentation; Collective intelligence; Human-agent Interaction

1. INTRODUCTION

The ability to process large amounts of data and forecast possible trends is fundamental in intelligence analysis, as well as in other analytic contexts such as web-commerce, social media, or criminal investigation. Current solutions for high-end analytics focus primarily on data aggregation and visualisation [7]. In this research, we focus on the process of making sense of the information extracted, or received from different sources. This more complex level of interpretation imposes a high cognitive burden on analysts since information may be unreliable, conflicting, or simply missing. Enabling collaborative analysis is one way to facilitate the review of evidence as it complements analysts' information [9]. More is required for collaborative analysis, however. Hypotheses inform strategies for preventing threats or coping with critical situations. To identify these, an analyst must combine several approaches to assess the available evidence, establish what information is credible, and understand what additional evidence may be required. This needs to be done in a timely manner to enable accurate interpretation of the available evidence which poses significant challenges for individual analysts. This combined reasoning process is resistant to automation due to the significant knowledge engineering effort required to process data and reasoning patterns [12]. The challenge we address here is: *how to develop agents to support this combined approach to reasoning with evidence throughout the process of analysis?*

We propose a model of reasoning with different kinds of evidence to be applied in collaborative settings for the review of hypotheses. Collaboration between analysts is enabled via *CISpaces*, an agent-based tool for constructing and sharing analyses [23]. CISpaces is focussed on facilitating the core phase of collaborative sensemaking within the intelligence analysis process, rather than presentation or information collection as in existing systems [3, 5, 24]. Our model of reasoning with diverse evidence is employed by CISpaces *interface agents* [21] to provide active support to analysts. *Argumentation schemes* [25], as patterns of defeasible inference, lie at the heart of our tool to structure inferences and evidence. To support analysts, agents employ argument schemes combined with: *argumentation* [18] to identify plausible hypotheses; *crowd-sourcing* [4, 11, 26] to form structured requests for information from groups of contributors and to enable semi-automated analysis of collected evidence; and *provenance* [8] to record meta-data about the origin of information and analyses, and to establish the credibility of hypotheses. Our contribution in this paper is to bring these approaches together to support the reasoning that underpins the complexity of the analytical process within our collaborative tool for intelligence analysis, CISpaces.

2. TOOL OVERVIEW AND SCENARIO

Here we briefly introduce our collaborative agent-based tool. CISpaces is developed according to the procedural phases of intelligence analysis defined in Pirolli and Card [17]. They model analysis as an iterative process of foraging for information to be collected and filtered from sources,

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

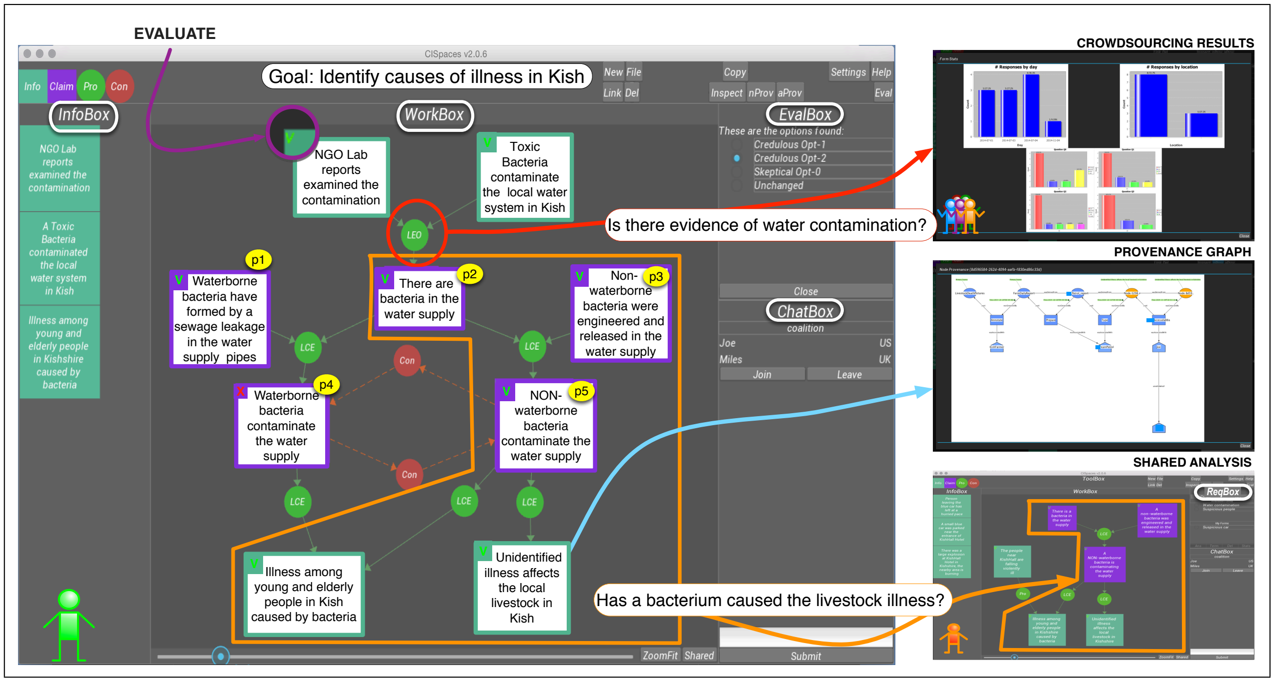

Figure 1: CISpaces interface and collaboration (Note that the nodes are overwritten for readability purposes)

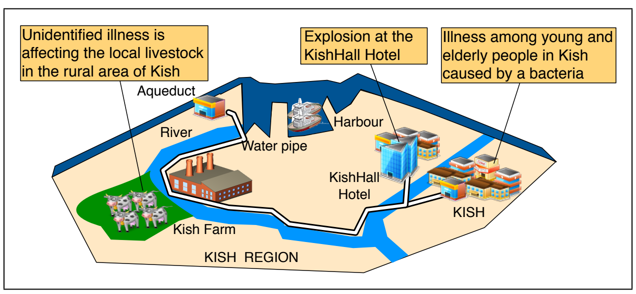

Figure 2: Initial information in Kish

and of interpreting this information by drawing inferences to identify hypotheses. CISpaces enables analysts to perform the core phases of analysis via the interface and a backend system that enables collaboration.

The CISpaces interface (Fig. 1) has two core components: the *InfoBox*, where collected information relevant to a task are streamed from external sources, typically from intelligence reports; and an individual *WorkBox*, the analytical space for analysts to construct hypotheses. In the *WorkBox* analysts import information, Info-nodes, or add new claims as Claim-nodes, represented as boxes. Each node has an attached provenance chain: data representing the phases of manipulation that a piece of information has gone through. The *WorkBox* is based upon a graphical representation of arguments [20, 24] where users can draw supporting "Pro-links" or defeating "Con-links" between nodes, represented by green or red circles respectively, forming arguments and attacks. Links can be annotated to provide additional meta-information about the type of inference between nodes.

Different forms of collaboration are supported in CISpaces. Analysts engage in a dialogue with other analysts via the *ChatBox* and a shared *WorkBox* permits the sharing of individually formed arguments. An analyst may also canvas groups of contributors via the *ReqBox*, by creating forms for collecting structured information using crowdsourcing.

We use a scenario throughout the paper that has been developed with the help of experienced professional intelligence analysts from different countries. Analysts are engaged in the investigation of water contamination in the fantasy city of Kish. Reports from the rural areas of Kish indicate an unidentified illness affecting local livestock and an increase in patient mortality. Intelligence analysts, from a coalition operating in the region, identify the contamination of drinking water as a possible cause. An intelligence requirement is issued to identify whether this is accidental, or related to other suspicious activities such as a local hotel explosion. CISpaces supports the team in analysing the situation as well as in liaising with local authorities to gather information about the spread of the illness in the region. Figure 2 shows an overview of the scenario.

3. REASONING WITH EVIDENCE

In CISpaces, analysts are assisted by *Interface Agents* [21] as those that employ agent-based techniques and interact with the users to provide support for tasks that they are attempting to execute. Interface agents collect relevant information from the user to start tasks, present further relevant information, and assist with complex problem solving. In this section, we describe how these agents combine argumentation-based reasoning [18, 25], reasoning over information provenance [22], and crowdsourcing [11] for information collection to assist analysts in reviewing evidence.

The procedural steps of analysis are underpinned by different evidential reasoning processes with three objectives:

- Identify what to believe happened from the claims constructed upon information (the *sensemaking process*);

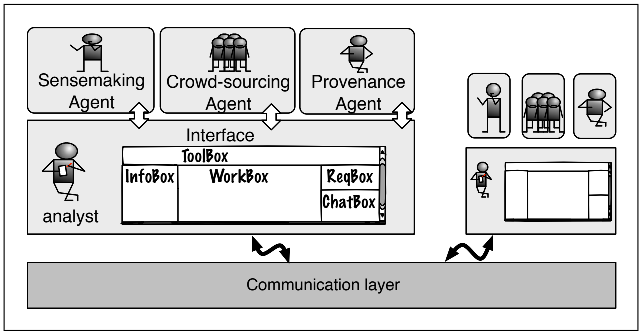

Figure 3: CISpaces Agent Support

- Derive conclusions from data aggregated from explicitly requested information (the *crowdsourcing process*);
- Assess what is credible according to the history of data manipulation (the *provenance reasoning process*).

In order to meet these objectives, the CISpaces interface enables analysts to access agent-support via:

- The *sensemaking agent* that employs argument schemes to guide critical review of evidence, and a model of argumentation to identify plausible hypotheses;
- The *crowd-sourcing agent* that interprets responses to structured requests for information from groups of collectors and feeds the results back into the analysis; and
- The *provenance agent* that inspects the provenance of information and identifies critical meta-data that may inform the credibility of the identified hypotheses.

In Figure 3 we show the architecture of CISpaces; note that analysts can interact with others but agents only interact with individual users. The exchange of arguments is based upon the Argument Interchange Format (AIF) [19], which may enable agent communication to provide more autonomous support for collaboration in future work.

4. ARGUMENTS FOR SENSEMAKING

The sensemaking process focusses on understanding what to believe about a situation given the information available. Different explanations of the same events may be possible, as information may be conflicting. CISpaces helps analysts perform sensemaking in collaboration, which permits different analysts to contribute different views on the problem. It also exploits argument mapping, which supports the structuring of analyses and the exchange of these different views. Given the typically large amount of information available, however, analysts may fail to identify weak points. Cognitive biases, such as confirmation (i.e., considering only information that confirms one's beliefs) may affect the accuracy of conclusions. Here, we discuss a model of argument schemes and argumentation for agents to support analysts to address these problems. Argument schemes are patterns of defeasible inferences [25], and guide the formulation of arguments in argumentation theory [6]. Computational instantiations are based on abstract atomic arguments and attacks (relations between arguments) and provide methods for deriving the acceptability status of arguments [6, 18].

4.1 Argumentation framework

In order to identify plausible conclusions, we adapt Prakken's argumentation framework [18], more recently developed into ASPIC+[13]. The sensemaking agent converts the WorkBox map of arguments composed of nodes, and Pro/Con-links built by the user, and evaluates this map according to argumentation semantics. The agent extracts conclusions that may be supported and displays them via the CISpaces interface. At present, we restrict the framework to ordinary premises and defeasible rules without preferences.

Definition 1. *An* argumentation system *AS is a tuple $\langle \mathcal{L}, \bar{\ }, \mathcal{R} \rangle$ where $\mathcal{L}$ is a logical language, $\bar{\ }$ is a contrariness function, and $\mathcal{R}$ is a set of defeasible rules. The contrariness function $\bar{\ }$ is defined from $\mathcal{L}$ to $2^{\mathcal{L}}$, s.t. given $\varphi \in \bar{\phi}$ with $\varphi, \phi \in \mathcal{L}$, if $\phi \notin \bar{\varphi}$, φ is called the* contrary *of ϕ, otherwise if $\phi \in \bar{\varphi}$ they are* contradictory *(including classical negation $\neg$). A* defeasible rule *is $\varphi_0, \ldots, \varphi_j \Rightarrow \varphi_n$ where $\varphi_i \in \mathcal{L}$.*

We refer to a rule $\alpha \Rightarrow \beta$ as r, where α is the *antecedent* and β is the *consequent*.

Definition 2. *A* knowledge-base *K in AS is a subset of the language $\mathcal{L}$. An* argumentation theory *is $AT = \langle K, AS \rangle$.*

An *argument Arg* is derived from the knowledge-base K of a theory AT. Let $Prem(Arg)$ indicate the premises of Arg, $Conc(Arg)$ the conclusion, and $Sub(Arg)$ the subarguments:

Definition 3. *An argument Arg is defined as:*
- $Arg = \{\varphi\}$ *with* $\varphi \in K$ *where* $Prem(Arg) = \{\varphi\}$, $Conc(Arg) = \varphi$, $Sub(Arg) = \{\varphi\}$.
- $Arg = \{Arg_1, \ldots, Arg_n \Rightarrow \phi\}$ *if there exists a defeasible rule in AS s.t.* $Conc(Arg_1), \ldots, Conc(Arg_n) \Rightarrow \phi \in \mathcal{R}$ *with* $Prem(Arg) = Prem(Arg_1) \cup \cdots \cup Prem(Arg_n)$, $Conc(Arg) = \phi$ *and* $Sub(Arg) = Sub(Arg_1) \cup \cdots \cup Sub(Arg_n) \cup Arg$.

Attacks are defined as those arguments that challenge others, defeats are those attacks that are successful:

Definition 4. *Given two arguments Arg_A and Arg_B:*
- Arg_A rebuts Arg_B *on* $Arg_{B'}$ *iff* $Conc(Arg_A) \in \bar{\varphi}$ *for* $Arg_{B'} \in Sub(Arg_B)$ *s.t.* $Arg_{B'} = \{Arg_{B1''}, \ldots, Arg_{Bn''} \Rightarrow \varphi\}$; *if* $Conc(Arg_A)$ *is contrary to* φ, Arg_A contrary-rebuts Arg_B.
- Arg_A undermines Arg_B *on* φ *iff* $Conc(Arg_A) \in \bar{\varphi}$ *such that* $\varphi \in Prem(Arg_B)$; *if* $Conc(Arg_A)$ *is contrary to* φ, Arg_A contrary-undermines Arg_B.

Definition 5. *An argument Arg_A* defeats *an argument Arg_B iff:* i) *Arg_A rebuts Arg_B on $Arg_{B'}$ and Arg_A contrary-rebuts $Arg_{B'}$; and* ii) *Arg_A undermines Arg_B on φ and Arg_A contrary-undermines Arg_B.*

An abstract argumentation framework AF corresponding to an AT includes a set of arguments as defined in Def. 3 and a set of defeats as in Def. 5. Sets of acceptable arguments (i.e., extensions) in an AF can be computed according to a semantics. The set of extensions that we consider here is $\hat{\xi} = \{\xi_1, ..., \xi_n\} \cup \{\xi_S\}$ such that each $\xi_i = \{Arg_a, Arg_b, \ldots\}$. The extensions $\xi_1, ..., \xi_n$ are the *credulous-preferred extensions* identified via preferred semantics; i.e., maximal w.r.t. set inclusion extensions that are conflict free (i.e., no arguments in any extension defeat each other), and admissible (i.e., each argument in the extension is defended against defeats from "outside" the extension). The *skeptical-preferred extension* ξ_S is the unique intersection of the credulous-preferred extensions. Table 1 shows an example of an AS and the related AT.

Table 1: WAT for partial analysis in Figure 1

$K = \{p_1, p_2, p_3\}$ $\ ^{-} = \{(p_4, p_5), (p_5, p_4)\}$	
$\mathcal{R} = \{p_1, p_2 \Rightarrow p_4; p_2, p_3 \Rightarrow p_5;\}$	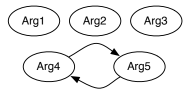
$Arg_1 : p_3;$ $Arg_2 : p_2;$ $Arg_3 : p_1;$	
$Arg_4 : Arg_1, Arg_2 \Rightarrow p_5;$	
$Arg_5 : Arg_2, Arg_3 \Rightarrow p_4;$	$(p_1, V)(p_2, V)(p_3, V) \in O_{1/2/S}$
$\xi_1 = \{Arg_1, Arg_2, Arg_3, Arg_5\}$	$(p_4, V), (p_5, X) \in O_1$
$\xi_2 = \{Arg_1, Arg_2, Arg_3, Arg_4\}$	$(p_4, X), (p_5, V) \in O_2$
$\xi_S = \{Arg_1, Arg_2, Arg_3\}$	$(p_4, ?), (p_5, ?) \in O_S$

4.1.1 The WorkBox Argumentation Theory

Here, we define the mapping of a WorkBox view to the corresponding AT, called WAT. An edge in CISpaces is represented textually as $\mapsto$, an info/claim node is written p_i and a link node is referred to as ℓ_{type} where $type = \{Pro, Con\}$. Then, $[p_1, \ldots, p_n \mapsto \ell_{Pro} \mapsto p_\phi]$ indicates that the Pro-link has $p_1, \ldots, p_n$ as incoming nodes and an outgoing node p_ϕ.

Definition 6. *A WAT is a tuple $\langle K, AS\rangle$ such that $AS = \langle \mathcal{L}, ^{-}, \mathcal{R}\rangle$ is constructed as follows:*
- *$\mathcal{L}$ is a propositional logic language, and a node corresponds to a proposition $p \in \mathcal{L}$. The WAT set of propositions is $\mathcal{L}_w$.*
- *The set $\mathcal{R}$ is formed by rules $r_i \in \mathcal{R}$ corresponding to Pro links between nodes such that: $[p_1, \ldots, p_n \mapsto \ell_{Pro} \mapsto p_\phi]$ is converted to $r_i : p_1, \ldots, p_n \Rightarrow p_\phi$*
- *The contrariness function between elements is defined as:* i) *if $[p_1 \mapsto \ell_{Con} \mapsto p_2]$ and $[p_2 \mapsto \ell_{Con} \mapsto p_1]$, p_1 and p_2 are contradictory;* ii) *$[p_1 \mapsto \ell_{Con} \mapsto p_2]$ and p_1 is the only premise of the Con link, then p_1 is a contrary of p_2;* iii) *if $[p_1, p_3 \mapsto \ell_{Con} \mapsto p_2]$ then a rule is added such that p_1 and p_3 form an argument with conclusion p_h against p_2, $r_i : p_1, p_3 \Rightarrow p_h$ and p_h is a contrary of p_2*[1]*.*

Definition 7. *K is composed of propositions p_i, $K = \{p_j, p_i, \ldots\}$, such that:* i) *let a set of rules $r_1, \ldots, r_n \in \mathcal{R}$ indicate a cycle such that for all p_i that are consequents of a rule r exists r' containing p_i as antecedent, then $p_i \in K$ if p_i is an info-node;* ii) *otherwise, $p_i \in K$ if p_i is not consequent of any rule $r \in \mathcal{R}$.*

4.1.2 Agent support for hypotheses identification

The sensemaking agent uses a WAT translation of a WorkBox to evaluate plausible conclusions and shows available options to the user as shown in Figure 1.

Definition 8. *Given an AF corresponding to a WAT, a proposition p_i and an existing extension ξ_j, p_i is acceptable if there is an argument $Arg_i \in \xi_j$ that has conclusion p_i.*

Given the set of all extensions $\hat{\xi}$ in the WAT, the analyst is presented with n colouring options that indicate when a node contains a statement that can be supported, unsupported or undecided. A node is supported if it contains a piece of information that is acceptable or defendable against its defeaters. A node is unsupported if it is rejected, and undecided if it has insufficient grounds to be either supported or unsupported.

Definition 9. *The set of options $\mathcal{O} = \{O_1, \ldots, O_n\}$ for a WAT is a set of cardinality $|\mathcal{O}| = |\hat{\xi}|$ where each option $O = \{(p_i, col_i)$ s.t. $p_i \in \mathcal{L}_w, col_i \in \{V, X, ?\}\}$. The assignment of col_i for p_i given an extension $\xi_j \in \hat{\xi}$ is: $col_i = V$ (supported), if p_i is acceptable in ξ_j; $col_j = X$ (unsupported), if p_i is a conclusion of an argument Arg_A that is defeated by $Arg_B \in \xi_j$; otherwise $col_j =?$ (undecided).*

We refer to the set of supported conclusions as the supported elements of an option O_i^V. Each of these options is proposed to the analyst. Table 1 represents a partial WAT for the analysis in Figure 1.

[1]This overcomes the syntactic limitation of [18] which does not allow to express that $p_1 \wedge p_3$ is a contrary of p_2.

4.2 Argumentation schemes

Here, we show a model of argumentation schemes that supports the structuring of analyses, and can be interpreted by sensemaking agents to drive further critical analysis.

Argumentation schemes are reasoning patterns that commonly occur in human reasoning and dialogue [25]. They represent templates for making presumptive inference, formed by premises supporting a conclusion and critical questions (CQs) that can be put forward against an argument. A commonly used example is the *argument from expert opinion*, used to describe an assertion warranted by expertise:

- Source E is an expert in domain S containing A,
- E asserts that proposition A is true,

$\Rightarrow$ Therefore, A may plausibly be true.

In a WAT each scheme corresponds to a rule. For example, the rule composing the expert opinion scheme is [18]:

$$r_{EO} : expert(E, A), assert(E, A), within(A, S), credible(E, S), reliable(E), evidence\_sup(A) \Rightarrow hold(A)$$

Critical questions include: *CQEO1* "How credible is E as an expert source?"; *CQEO2* "Is E an expert in the field that A is in?"; *CQEO3* "Is A consistent with the testimony of other experts?". The critical questions are used as pointers to other arguments that may challenge this inference. CQs are mapped to a WAT according to the type of attack as:

- *undermining CQs* as attacks to premises: a Con-link is mapped to a contradictory relation. For example, CQEO2: $\neg expert(E, A)$
- *undercutting CQs* challenging an exception of the inference rule. In CISpaces, undercuts are contrary underminers to propositions implicit in the scheme; e.g. CQEO1: $non\_credible(E, S)$ contrary of $credible(E, S)$
- *rebutting CQs* as contradictions of the conclusions: a Con-link is mapped to a contradictory relation. For example, CQEO5: $\neg hold(A)$.

Agents use critical questions to help users identify weak points in analyses. For this, we model the type of inferences that represent the sensemaking process using some of the argumentation schemes from Walton et al. [25]. Here, we show how these schemes can be linked together to build hypothetical explanations for situations.

4.2.1 Representing intelligence

Intelligence analysis broadly consists of three components: **Activities** (Act) including *actions* performed by actors, and *events* happening in the world; **Entities** (Et) including actors as individuals or groups, and objects such as resources; and **Facts** (Ft) including statements about the state of the world regarding entities and activities. Important relations between these elements include: *causal* relations representing the distribution of activities, their correlation and causality; and relations that connect entities and activities through temporal, geographic or thematic *association*. Heuer [9] argues that analysts tend to adopt a *historian approach*, making use of these relations to reconstruct a narrative that explains events. This is also advocated by recent work on schemes for criminal investigation [2], where the evidence is used to prove the plausibility of a story already prepared. In intelligence analysis, additional arguments and critical questions are needed to link together information to construct hypothetical stories. Intelligence elements act as premises for inferences, and conclusions are tentatively drawn by discovering relations between these elements. According to the type of relation (causal or associative) we can now instantiate two main types of schemes for the sensemaking process.

4.2.2 Causal Relations

A *hypothesis in intelligence analysis* is composed of activities and events that show how the situation has evolved. The *argument from cause to effect* forms the basis of these hypotheses. This is referred to as Arg_{CE}, and considers a cause $\mathcal{C}$, a fact Ft_i or an activity Act_i, its effect $\mathcal{E}$, also a fact or an activity, and a causal rule that links $\mathcal{C}$ to $\mathcal{E}$. The scheme, adapted from [25], is:

- Typically, if $\mathcal{C}$ occurs, then $\mathcal{E}$ will occur
- In this case, $\mathcal{C}$ occurs

$\Rightarrow$ Therefore, in this case $\mathcal{E}$ will occur

This argument can, for example, be used to state that a waterborne bacterium may cause the illness among people in Kish. Formally Arg_{CE} is:

$$r_{CE} : \mathit{rule}(R, \mathcal{C}, \mathcal{E}), \mathit{occur}(\mathcal{C}), \mathit{before}(\mathcal{C}, \mathcal{E}), \mathit{ruletype}(R, \mathit{causal}), \mathit{noexceptions}(R) \Rightarrow \mathit{occur}(\mathcal{E})$$

The critical questions are:

- CQCE1: Is there evidence for $\mathcal{C}$ to occur?
- CQCE2: Is there a general rule for $\mathcal{C}$ causing $\mathcal{E}$?
- CQCE3: Is the relationship between $\mathcal{C}$ and $\mathcal{E}$ causal?
- CQCE4: Are there any exceptions to the causal rule that prevent the effect $\mathcal{E}$ from occurring?
- CQCE5: Has $\mathcal{C}$ happened before $\mathcal{E}$?
- CQCE6: Is there any other $\mathcal{C}'$ that caused $\mathcal{E}$?

A set of arguments of type Arg_{CE}, where each effect $\mathcal{E}$ is a premise $\mathcal{C}$ for the next Arg_{CE}, constitutes a hypothesis that may then be reviewed through critical questions. Question CQCE4 leads to other rebutting arguments of type Arg_{CE} with conclusion $\neg \mathit{occur}(E)$. CQCE5 is an undercut on $\mathit{before}(\mathcal{C}, \mathcal{E})$ that distinguishes the type of reasoning required to piece together information for a story and the counterpart of identifying evidence for a story as in [2]. The questions CQCE1 and CQCE2 point towards undermining arguments with conclusions $\neg \mathit{occur}(\mathcal{C})$ and $\neg \mathit{rule}(R, \mathcal{C}, \mathcal{E})$, while CQCE3 identifies an undercut on $\mathit{ruletype}(R, \mathit{causal})$. These CQs are used to challenge whether $\mathcal{C}$ has occurred and an analyst may defend the argument using: an *argument from analogy* reporting a case with similarities to $\mathcal{C}$; an *argument from rule* debating a rule that leads to the current situation $\mathcal{C}$; or an *argument from sign* to explain that $\mathcal{C}$ is likely to happen if its sign is verified [25].

Critical question CQCE6 has a different purpose. In CISpaces, analysts are required to represent a cause as a Pro link to an effect. However, analysts may have evidence for the effect and infer a plausible cause using abductive reasoning and, in this case, alternative causes must be considered. CQCE6 is used to consider these alternatives, and the role of agents is to interpret the critical question as a rebuttal for $\mathcal{C}$ within an *abductive argument from effect to cause* [25]. CQCE6 results in alternative incoming nodes to the Pro-link representing a contradictory relation between causes.

4.2.3 Associative Relations

The sensemaking process may shift from understanding what happened to understanding what entities were involved and their association to the activity. We use Arg_{ID}, an *argument for identifying an agent from past actions* [25]:

- An activity Act_i occurs, and Et_i may be involved
- To bring about activity Act_i some property H is required
- Et_i fits the property H

$\Rightarrow$ Therefore, Et_i is associated with Act_i

Properties H are facts Ft of type "Et_i is affected by Act_i/Et_j" or "Et_i is in the location Et_j of Act_i". Formally, Arg_{ID} is:

$$r_{ID} : \mathit{occur}(Act_i), \mathit{requiresProp}(Act_i, H), \mathit{entity}(Et_i), \mathit{hasProp}(Et_i, H), \mathit{noexceptions}(Et_i, H, Act_i) \Rightarrow \mathit{association}(Et_i, Act_i)$$

And the critical questions are:

- CQID1: Has the activity Act_i happened?
- CQID2: Does the activity Act_i fit the properties listed?
- CQID3: Does Et_i fit the properties required by Act_i?
- CQID4: Are there other entities that fit properties H?
- CQID5: Is there an exception to property H that undermines the association between entity and activity?

Scheme Arg_{ID} can be used to assert that activity Act_i is "an unidentified bacterium has contaminated the water supply" and we conclude that "E.coli bacterium has contaminated the water supply". CQID1, CQID2 and CQID3 are underminers for the argument with conclusions $\neg \mathit{occur}(Act_i)$, $\neg \mathit{hasProp}(Et_i, H)$, or $\neg \mathit{requiresProp}(Act_i, H)$. CQID4 is a rebutting argument concluding $\neg \mathit{association}(Et_i, Act_i)$ and CQID5 is an undercut on $\mathit{noexceptions}(Et_i, H, Act_i)$. CQID3 focusses on properties about the entity associated with the activity. However, this association may be derived by other associative relations. Analysts use matrices to associate people to people, or groups, and people to locations, resources to people, and so on [9]. To reply to CQID3 we may use: *arguments from transitivity* where a presumptive transitivity relationship is applied; *arguments from the group* where properties of a member are applied to the group; or *arguments from verbal classification* where properties of a group are applied to the members [25].

The link between the major schemes for causal and associative relations is identified by question CQCE, stating that some entity Et_i was associated with the cause. Failing to provide evidence for this will invalidate the conclusion that an effect may happen. This CQ is a link from causal relations to an associative argument Arg_{ID}. On the other hand, Arg_{CE} may answer question CQID1 concerning whether the activity happened, linking association to causality.

4.2.4 Agent support for inference-making

In CISpaces, analysts build individual and collaborative maps of arguments by linking nodes through Pros and Cons. When the users add a Pro-link between nodes, they may annotate this link with reasons for how conclusions follow from the premises. For example, a label may be "Cause to Effect" if the Pro inference is a causal relation, $[p_1, p_3 \mapsto \ell_{Pro(LCE)} \mapsto p_2]$. The sensemaking agent interprets the annotated link as one of the schemes and suggests critical questions. Agents and analysts review the analysis as follows.

Step 1. Analysts annotate each incoming node $p_1, \ldots, p_n$

of the Pro-link as premises of the corresponding scheme; e.g. the premise "there are waterborne bacteria that contaminate the water" is identified with "In this case, $\mathcal{C}$ occurs".

Step 2. The critical questions of a matched scheme are employed by the agent to drive further analysis of the topics. The agent shows only relevant critical questions about the premises; e.g., "CQCE1: Is there evidence for $\mathcal{C}$ to occur?".

Step 3. Analysts may then select the questions that must be answered in order for the conclusion to be acceptable.

Step 4. The agent generates a negative answer in a new node connected via a Con-link; e.g., "There is no evidence for $\mathcal{C}$ to occur". If the answer is not counterattacked, the conclusions of the argument may not be supported.

By employing argumentation schemes agents guide sense-making, and help analysts identify weak points in the analysis. This approach extends existing argument mapping systems, e.g. [20, 24], with active support from agents that propose critical questions to drive further analysis. Given this structure of arguments, a WAT is more than a set of inferences: it represents the causal transition between events and associated entities, and alternative options. When evaluating a WAT, the supported elements of option O_i^V, suggested by the agent, represents a plausible hypothesis.

5. CROWD-SOURCED EVIDENCE

In this section, we describe the crowdsourcing process to extract evidence from structured requests. During analysis, further information may be needed to draw conclusions on plausible hypotheses. Analysts may not have sufficient resources to obtain this information, but this may be met by asking a crowd of collectors. Crowdsourcing is a technique that uses human computation to sense information and discover truth in a timely, large-scale and cost-efficient manner [4, 11, 26]; this is particularly effective in event detection [15]. Our approach is distinct from more traditional crowdsourcing because we explore how agents can interpret results and introduce them into the analysis in a meaningful way. The crowdsourcing agent assists the analyst by employing principled aggregation algorithms combined with argumentation schemes.

Step 1) Initialise a task. In CISpaces, a new task starts by selecting a CQ for the crowd or a node p_t that is interpreted by the agent as "Is there evidence for p_t?". For simplicity, we only consider boolean (yes/no) tasks.

DEFINITION 10. *A crowdsourcing task T is a tuple $T = \langle p_t, q_t, d_t, n_t, c_t, Q\rangle$ where q_t is the task definition, p_t is a proposition in $\mathcal{L}$ representing the answer to q_t, $Q = \{q_1, \ldots, q_n\}$ is a set of questions relevant to p_t, d_t is the deadline for the task, n_t the number of participants, and c_t the target crowd.*

An analyst, for example, may define a location-sensitive task T to probe the water contamination. Participants are those living in Kish, and they have two days to report colour and temperature of the tap water. Task T is:
- p_t: "The water in Kish is contaminated"
- q_t: "Is there evidence that the water is contaminated?"
- c_t: People that live in the region of Kish
- d_t: starting day+2 days, n_t: 10, $Q = \{q_1, q_2\}$

Step 2) Define forms. When the task is initiated, the analyst creates a form with n questions q_k to be answered by the contributors. The form requires additional information that helps agents draw conclusions from the collected results.

DEFINITION 11. *A question $q_i \in Q$ is a tuple $q_i = \langle type_i, text_i, options_i\rangle$, where $type_i$ is either* categorical *or* numerical*, $text_i$ defines the question asked to the crowd and $options_i$ indicate the space of possible answers.*

Let ev be a function that maps a number/category to its evaluation $\{Pro, Con\}$, the types of questions are:

- $type_i = categorical$: $Cat = \{cat_1, \ldots, cat_n\}$ is the space of possible answers for q_k; for each category cat_j the analyst chooses $ev(cat_j) = Pro$ (or $ev(cat_j) = Con$) if cat_j is a reason for believing $p_t \in T$(or $\neg p_t$).
- $type_i = numerical$: the answers are real numbers $n \in \mathbb{R}$; analysts define $ev(n) = \{Pro, Con\}$ as specific values for n to be considered as Pro or Con for p_t.

In our example, the questions q_1, q_2 can be defined as:
- $q_1 = \langle categorical$,"What colour is your tap water?",$ev(cat)$ $=\{(Clear, Con), (Brown, Pro), (Yellow, Pro), (White, Con)\}\rangle$
- $q_2 = \langle numerical$,"What is the temperature of your cold water?",$ev(n) = \{(n < 20, ev = Con), (n \geq 20, ev = Pro)\}\rangle$

Step 3) Collect reports. The crowdsourcing agent sends the query to the appropriate type of crowd c_t, based on the details of the registered collectors. The task terminates when it reaches the deadline d_t or the number of reports n_t are acquired. A report $\hat{\Omega}^j$ for participant j contains an answer $\hat{\omega}_k^j$ for each question q_k, $\hat{\Omega}^j = \{\hat{\omega}_1^j, \ldots, \hat{\omega}_m^j\}$. We only consider valid and complete reports.

For an m-category question q_k, let x_i be the number of participants that reported cat_i s.t. $\hat{\omega}_k^j = cat_i$, the vector $\bar{x} = (x_1, \ldots, x_m)$ represents the count for s participants.

For a numerical q_k the report is a number $\hat{\omega}_k^j = y_i$. A set $Y_k = \{y_1, \ldots, y_s\}$ represents the reports for s participants.

Step 4) Aggregation of results. For categorical data we are interested in knowing the probability of the categories of a multi-valued answer to question q_k. We use a Dirichlet distribution that captures the probability of the m possible outcomes (corresponding to the m categories) in an m-component random probability variable $\bar{\pi} = (\pi_1, \ldots, \pi_m)$, and $\pi_i \geq 0, \sum_{i=1}^m \pi_i = 1$. A Dirichlet distribution with priors [10] takes into consideration an initial belief prior to obtaining evidence for the situation. The analysis of data is then a posterior Dirichlet distribution that combines prior beliefs and collected reports for question q_k. The distribution is defined on m mutually disjoint categories Cat, the vector $\bar{x} = (x_1, \ldots, x_m)$ is interpreted as the evidence vector of collected responses, a vector $\bar{\alpha} = (\alpha_1, \ldots, \alpha_m)$ represents the base rate over the same elements, $\alpha_i \geq 0, \sum_{i=1}^m \alpha_i = 1$.

DEFINITION 12. *The Dirichlet distribution with priors that interprets the categorical data for a question q_k is:*

$$f(\bar{\pi}, \bar{x}, \bar{\alpha}) = \frac{\Gamma(\sum_{i=1}^m x_i + C\alpha_i)}{\sum_{i=1}^m \Gamma(x_i + C\alpha_i)} \Pi_{i=1}^m \pi_i^{x_i + C\alpha_i - 1}$$

The expectation of π_i that the answer to q_k is cat_i, is:

$$E[\pi_i|\bar{x}, \bar{\alpha}] = \frac{(x_i + C\alpha_i)}{C + \sum_{i=1}^m x_i}$$

We use $C = 2$ as *a priori* constant assuming a uniform distribution over the answers (as is often used in the literature [10]) and α_i as default base rate $1/m$. The vector $\bar{\epsilon}_k = (E[\pi_1], \ldots, E[\pi_m])$ refers to the resulting expected values for the m categories of question q_k. A more sophisticated approach would consider crowd features such as reliability and location by manipulating the prior.

For numerical data, we consider a weighted mean of the s collected reports Y_k for q_k. In the simplest case weights w_i are assumed to be 1, although these may vary according to features of the reports as for the prior probability.

DEFINITION 13. *The interpretation of a numerical q_k is:*

$$\mu_k = \frac{\sum_{i=1}^{s} w_i y_i}{\sum_{i=1}^{s} w_i}$$

Step 5) Analysis of results. After aggregating the results for each question q_k, the crowdsourcing agent uses the task definition T to *automatically* build a partial argument map that is integrated within the overall analysis. The *argument from generally accepted opinion* [25], Arg_{CS}, represents the defeasible inference that a statement is plausible if a significant majority in a group accepts it. Critical questions focus on whether the crowd is believable, or corroborating evidence is needed to accept the conclusions. Arg_{CS} is:

- Given that the crowd was asked $text_k$ and (p_{cs1})
- Answer A is generally accepted as true (p_{cs2})

⇒ Therefore, A may plausibly be true (p_{cscn})

The crowdsourcing agent constructs an argument Arg_{CS} for each question q_k: $text_k$ is the text defined for q_k and answer A corresponds to the mean μ_k for numerical questions, or cat_j^k for categorical ones where cat_j^k is the category with maximal expected value ϵ_j in $\bar{\epsilon}_k$. In CISpaces, Arg_{CS} is a Pro-link $[p_{cs1}, p_{cs2} \mapsto \ell_{Pro(LCS)} \mapsto p_{cscn}]$. Formally:

r_{CS} : $request(K, C), crowd(C), assert(C, A),$
$credible(C, A), reliable(C), evidence\_sup(A) \Rightarrow hold(A)$

Argument Arg_{CS} simply reports outcomes from crowdsourcing tasks, but agents can provide additional support in analysing the meaning of such results. The boolean proposition p_t represents the crowdsourcing task, and each answer for q_k is defined with options that can be used to interpret the results. If the category is associated with a Pro, $ev(cat_j^k) = Pro$ the agent creates a new link such that $[p_{cscn} \mapsto \ell_{Pro} \mapsto p_t]$. For the mean, the link is a Pro if $ev(\mu_j) = Pro$. Otherwise the link created is a Con. If two or more categories have the maximal expectation, a number of Arg_{CS} will lead to the final Pro/Con-link for p_t.

Assume we have collected 10 reports for q_1, q_2 such that:

- q_1: $Cat_1 = \{Clear, Brown, Yellow, White\}$, $\bar{x}_1 = (6, 1, 2, 1)$ with expectation $\bar{\epsilon}_1 = (0.542, 0.125, 0.208, 0.125)$
- q_2: $Y_2 = \{21, 22, 25, 24, 18, 17, 22, 20, 23, 19\}$ s.t. $\mu_2 = 21$

Two links are automatically generated in the WorkBox following Arg_{CS}, with conclusions (Arg_{CS1}) "The colour of the tap water is Clear" and (Arg_{CS2}) "The temperature of the tap water is 21". Given $ev(Clear) = Con$ the conclusion of Arg_{CS1} is a Con for the water being contaminated (p_t). For $ev(\mu_2) = Pro$ the premises of Arg_{CS2} indicate that the water has a higher temperature, hence, there is evidence for the contamination and p_{cscn} will be linked to p_t with a Pro-link.

Gathering additional information is necessary to avoid the rejection of hypotheses on the basis of insufficient evidence [9]. Our approach to crowdsourcing, evidence interpretation and *automated* integration of the outcome(s) into an analysis provides an effective method integrate this form of human intelligence into the sensemaking process.

6. REASONING ABOUT PROVENANCE

The origins of information (including information from the crowd), and how and by whom this information is interpreted during analysis are important to establish the credibility of hypotheses. Provenance can be used to annotate how, where, when and by whom some information was produced [14]. CISpaces enables analysts to inspect these provenance records. Consulting the provenance of each report, however, increases the cognitive load on the analyst, and our aim is to support effective sensemaking while managing this load. Agents support analysts by extracting important elements of provenance to introduce them into the sensemaking process via an argument scheme [22]. The scheme uses provenance to determine the credibility of information and relevant hypotheses. We explain here how this scheme complements the review of hypotheses, highlighting evidence provided by the provenance of the analysis.

Provenance is recorded in CISpaces as RDF triplets using the W3C standard PROV Data Model [14]. PROV-DM expresses provenance in terms of p-entities (A_{pv}), p-activities (P_{pv}), and p-agents (Ag_{pv}) that have caused an entity to be, and defines different relationships between these elements. We refer to p-elements to distinguish them from intelligence analysis elements. Figure 4 shows an example provenance graph for lab tests conducted on a contaminated water sample. The provenance chain of a node p_j is represented as a directed acyclic graph $G_P(p_j)$ of relationships between A_{pv}, P_{pv}, and Ag_{pv}. $G_P(p_j)$ is a joint path from the node containing p_j to its primary sources; i.e., sources that first produced the information. A provenance chain $G_P(p_j)$ can be queried as a graph pattern P_m which is a structured graph with nodes being variables on the p-elements. For example, the pattern $P_g(?a1, ?a2, ?p, ?ag)$ highlighted in Fig. 4 represents the generation of a p-entity $?a1$ derived from a p-entity $?a2$ by a p-activity $?p$, which was associated with a p-agent $?ag$. These patterns represent relevant provenance information that may warrant the credibility of p_j, and they can be integrated into the analysis by applying the *argument scheme for provenance* [22] (Arg_{PV}):

- Given p_j about activity Act_i, entity Et_i, or fact Ft_i (p_{pv1})
- $G_P(p_j)$ includes pattern P'_m of p-entities A_{pv}, p-activities P_{pv}, p-agents Ag_{pv} involved in producing p_j (p_{pv2})
- $G_P(p_j)$ infers that information p_j is true (p_{pv3})

⇒ Then, $Act_i/Et_i/Ft_i$ in p_j may plausibly be true (p_{pvcn})

Critical questions for this scheme are:

- CQPV1: Is p_j consistent with other information?
- CQPV2: Is p_j supported by evidence?
- CQPV3: Does $G_P(p_j)$ contain p-elements that lead us not to believe p_j?
- CQPV4: Is there any other p-element that should have been included in $G_P(p_j)$ to infer that p_j is credible?

A question "Can it be shown that the information is verifiable?" (e.g. CQID1, CQCE1) shifts the reasoning process to provenance. Questions CQPV1 and CQPV2 shift back to sensemaking by requiring further evidence for Act_i, Et_i, or Ft_i to be supported. To integrate the provenance elements into the analysis, the agent extracts and shows relevant patterns P_m to the analyst. The analyst, then, has two choices. If P'_m is considered as being relevant for asserting that p_j is plausible, the agent constructs a link $[p_{pv1}, p_{pv2}, p_{pv2} \mapsto \ell_{Pro(LPV)} \mapsto p_{pvcn}]$ in the WorkBox. The conclusion p_{pvcn} may already exist in the WorkBox since it concerns an Info or a Claim node, and so this link provides additional evidence for p_{pvcn}. However, a pattern P_m may be a reason for believing that p_j is not credible, based upon reasons expressed by CQPV3 or CQPV4. In this case, the

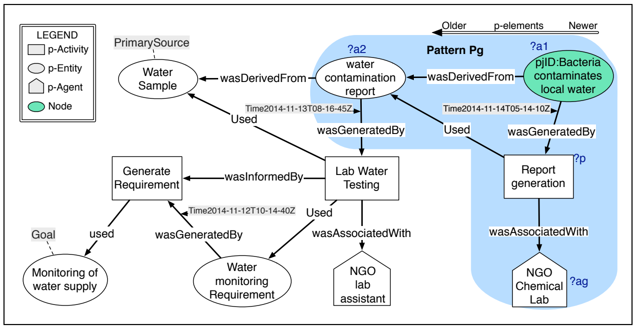

Figure 4: Provenance of water contamination tests

analyst must select the pattern. The agent builds a Pro-link with conclusion p_{pvcn}, and constructs a Con-link with a negative answer to the selected critical question, and conclusion p_{pv3}. This indicates an attack on the premises of Arg_{PV}, and so the conclusion would not be supported. Agents are able to support an analyst in extracting relevant provenance information to be consumed in the process of reviewing the credibility of evidence and hypotheses.

7. DISCUSSION

In this paper, we showed how agents support analysts in the identification of plausible hypotheses within CISpaces. We now illustrate this using our running example developed with professional analysts. Our initial analysis concerned an illness among both people and livestock in Kish (Section 2). In Section 4 we showed how the sensemaking agent helps an analyst structure and identify different hypotheses. Assume that there are two plausible causes for this illness: waterborne-bacteria or engineered non-waterborne bacteria contaminating the water supply. The crowdsourcing agent reports evidence of elevated water temperature (Section 5). However, the normal colour of the water may lead the analyst to conclude that there is no contamination, and so the causes of the illness may lie elsewhere. Nevertheless, the analyst may suggest that clear water does not exclude the presence of bacteria, particularly if artificially engineered, which reinstates the plausibility of water contamination. Moreover, the analyst intends to investigate how the waterborne bacteria hypothesis was formed, seeking support from the provenance agent. The agent brings forward a new argument (Section 6) indicating that the existence of waterborne bacteria in the water supply was stated by a group of biologists. While investigating the diffusion of aerial bacteria, the group gathered information about water diseases, although this was not their primary objective. Since the waterborne bacteria hypothesis results less credible (uncovered by the agent with CQPV3), the analyst may now conclude that the bacteriological outbreak is due to engineered bacteria. In order to deliver support, CISpaces agents employ argument schemes to structure the sensemaking, the analysis of crowdsourced results, and the inspection of provenance.

Agent-based argumentation is an established technique for dealing with conflicting evidence. Formal models are used to capture different types of conflicts arising between information [1, 25], to resolve these conflicts [18], and to evaluate the reliability of conclusions [16, 22]. Argumentation techniques, however, focus on agent-based decision-making, and such methods may require training to be used by analysts due to the extensive formalisation required. On the other hand, argument mapping provides intuitive and effective support for critical thinking [20, 24], but does not offer agent support for reasoning. Here, we combine these approaches to enable analysts to directly interact and benefit from a computational model of argumentation in the construction and evaluation of hypotheses.

In order to support analysts to better select hypotheses, we employ crowdsourcing to facilitate the acquisition of additional evidence and provenance to explore the credibility of information. In recent research, agent-based approaches have been applied to crowdsourcing to automate decision-making on behalf of the requestors such as whom to hire, [11], which is more akin to a trust decision making problem. More traditional approaches focus on result aggregation to mitigate biases from unreliable sources [4, 15, 26]. Similarly, work on provenance is primarily concerned with data quality and interoperability [8]. In this research, we study how agents can *interpret* provenance and crowdsourced data to assist analysts and *integrate* this information in generating coherent explanations of observed evidence.

8. CONCLUSION

In this paper, we have presented a combined approach to interpreting available evidence, assessing its credibility and identifying additional information requirements to construct hypotheses that are plausible explanations of a situation. Our primary contribution is the combination of different agent-based techniques for evidential reasoning and information gathering to support analysts in the delivery of more effective intelligence products. This is realised in CISpaces that employs state-of-the-art argumentation techniques to compute extensions, and crowdsourcing and provenance to model and support the reasoning that underpins the elaboration of data in complex analysis.

CISpaces has been developed and tested with the help of professional intelligence analysts. In addition to offering effective support for their work, analysts have highlighted that CISpaces is useful for training and provides an effective means to record an audit trail that includes important elements of the reasoning processes involved in the analysis of competing hypotheses. This audit trail provides a record of what evidence was available at the time of analysis, what alternative hypotheses were considered and which were considered most credible, which may be used to improve future analyses. Future evaluation of this research will require controlled experiments to be conducted with human subjects to quantify the benefits of CISpaces identified through qualitative feedback from professional analysts.

Acknowledgments

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

[1] F. Bex, H. Prakken, C. Reed, and D. Walton. Towards a formal account of reasoning about evidence: argumentation schemes and generalisations. *Artificial Intelligence and Law*, 11(2-3):125–165, 2003.

[2] F. Bex and B. Verheij. Solving a murder case by asking critical questions: An approach to fact-finding in terms of argumentation and story schemes. *Argumentation*, 26(3):325–353, 2012.

[3] D. Billman, G. Convertino, J. Shrager, P. Pirolli, and J. Massar. Collaborative intelligence analysis with CACHE and its effects on information gathering and cognitive bias. In *Human Computer Interaction Consortium*, Boulder, CO, 2006.

[4] D. C. Brabham. Crowdsourcing as a model for problem solving: An introduction and cases. *Convergence: The International Journal of Research into New Media Technologies*, 14(1):75–90, 2008.

[5] A. De Liddo, Á. Sándor, and S. Buckingham-Shum. Contested collective intelligence: Rationale, technologies, and a human-machine annotation study. *Computer Supported Cooperative Work*, 21:417–448, 2012.

[6] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[7] D. Fisher, R. DeLine, M. Czerwinski, and S. Drucker. Interactions with big data analytics. *interactions*, 19(3):50–59, 2012.

[8] O. Hartig and J. Zhao. Using web data provenance for quality assessment. *Proceedings of the First International Workshop on the Role of Semantic Web in Provenance Management at the International Semantic Web Conference*, 2009.

[9] R. J. Heuer. *Psychology of intelligence analysis*. US Government Printing Office, 1999.

[10] A. Jøsang and J. Haller. Dirichlet reputation systems. In *Proceedings of the Second IEEE International Conference on Availability, Reliability and Security.*, pages 112–119, 2007.

[11] E. Kamar, S. Hacker, and E. Horvitz. Combining human and machine intelligence in large-scale crowdsourcing. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems*, pages 467–474, 2012.

[12] J. Llinas. Challenges in information fusion technology capabilities for modern intelligence and security problems. In *Proceedings of the IEEE European Conference on Intelligence and Security Informatics Conference*, pages 89–95, 2013.

[13] S. Modgil and H. Prakken. The ASPIC+ framework for structured argumentation: A tutorial. *Argument & Computation*, 5(1):31–62, 2014.

[14] L. Moreau and P. Missier. PROV-DM: The PROV Data Model. Available at `http://www.w3.org/TR/prov-dm/`, April 2013.

[15] R. W. Ouyang, M. Srivastava, A. Toniolo, and T. J. Norman. Truth discovery in crowdsourced detection of spatial events. In *Proceedings of the Twenty-Third International Conference on Information and Knowledge Management*, pages 461–470, 2014.

[16] S. Parsons, Y. Tang, E. Sklar, P. McBurney, and K. Cai. Argumentation-based reasoning in agents with varying degrees of trust. In *Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems*, pages 879–886, 2011.

[17] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of the International Conference on Intelligence Analysis*, 2005.

[18] H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2010.

[19] I. Rahwan, F. Zablith, and C. Reed. Laying the foundations for a world wide argument web. *Artificial Intelligence*, 171(10–15):897 – 921, 2007.

[20] C. Reed and G. Rowe. Araucaria: Software for argument analysis, diagramming and representation. *International Journal on Artificial Intelligence Tools*, 13(04):961–979, 2004.

[21] K. Sycara, A. Pannu, M. Williamson, D. Zeng, and K. Decker. Distributed intelligent agents. *IEEE Expert*, 11(6):36–46, 1996.

[22] A. Toniolo, F. Cerutti, N. Oren, T. J. Norman, and K. Sycara. Making informed decisions with provenance and argumentation schemes. In *Proceedings of the Eleventh International Workshop on Argumentation in Multi-Agent Systems*, 2014.

[23] A. Toniolo, T. Dropps, R. W. Ouyang, J. A. Allen, T. J. Norman, N. Oren, M. Srivastava, and P. Sullivan. Argumentation-based collaborative intelligence analysis in CISpaces. In S. Parsons, N. Oren, C. Reed, and F. Cerutti, editors, *Computational Models of Argument*, volume 266 of *Frontiers in Artificial Intelligence and Applications*, pages 481–482. IOS Press, 2014.

[24] T. van Gelder. The rationale for Rationale$^{\mathrm{TM}}$. *Law, Probability and Risk*, 6(1-4):23–42, 2007.

[25] D. Walton, C. Reed, and F. Macagno. *Argumentation schemes*. Cambridge University Press, 2008.

[26] J. Whitehill, P. Ruvolo, T. Wu, J. Bergsma, and J. Movellan. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *In Advances in Neural Information Processing Systems 22*, pages 2035–2043. Curran Associates, Inc., 2009.